A Geometric Distance Oracle for Large Real-World Graphs

Deepak Ajwani*, W. Sean Kennedy†, Alessandra Sala*, Iraj Saniee†

April 22, 2014

arXiv:1404.5002v1 [cs.SI] 20 Apr 2014

Abstract

Many graph processing algorithms require determination of shortest-path distances between arbitrary numbers of node pairs. Since computation of exact distances between all node-pairs of a large graph, e.g., 10M nodes and up, is prohibitively expensive both in computational time and storage space, distance approximation is often used in place of exact computation. A *distance oracle* is a data structure that answers inter-point distance queries more efficiently than in standard $O(n^2)$ in time or storage space for an n node graph, e.g., in $O(n \log n)$. In this paper, we present a novel and scalable distance oracle that leverages the hyperbolic core of real-world large graphs for fast and scalable distance approximation via spanning trees. We show empirically that the proposed oracle significantly outperforms prior oracles on a random set of test cases drawn from public domain graph libraries. There are two sets of prior work against which we benchmark our approach. The first set, which often outperforms all other oracles, employs embedding of the graph into low dimensional Euclidean spaces with carefully constructed hyperbolic distances, but provides no guarantees on the distance estimation error. The second set leverages Gromov-type tree contraction of the graph with the additive error guaranteed not to exceed $2\delta \log n$, where δ is the hyperbolic constant of the graph. We show that our proposed oracle 1) is significantly faster than those oracles that use hyperbolic embedding (first set) with similar approximation error and, perhaps surprisingly, 2) exhibits substantially lower average estimation error compared to Gromov-like tree contractions (second set). We substantiate our claims through numerical computations on a collection of a dozen real world networks and synthetic test cases from multiple domains, ranging in size from 10s of thousand to 10s of millions of nodes.

1 Introduction

The explosion of available information in the past decade, in part due to the rapid shift towards online media and interactions has led many research, business and marketing communities to store and analyze very large data sets. Mining these data sets promises to reveal a wealth of information about the interests of and the kind of interactions between subscribers, groups, people, objects and even ideas. These interactions are often naturally represented by graphs, where for example, nodes correspond to the entities of interest and the (weighted) links represent the strength of the interaction between them. Graphs extracted for data mining are often massive, comprising of millions to billions of connections. At this scale, graph algorithms requiring $\Omega(n^2)$ computational steps or storage reach their useful limit in terms of run time and memory requirements. There is clearly a need for implementations of graph computational primitives at this scale.

Computing shortest path distance between arbitrary nodes of a graph is among such fundamental computational primitives. Many data mining schemes invoke this computational primitive in the scale of the number of nodes, and therefore it is imperative that this computation can be carried out very rapidly and with limited memory consumption. Distance oracles are among the many approaches that have been proposed and used for simplification of shortest distance computation for large-scale graphs. A distance oracle involves an auxiliary data structure which is cheaper to compute and fast to query. It should ideally satisfy the following four properties:

1. (Initial Processing Speed) the computation involved in the creation of the auxiliary data structure should be scalable, e.g. be $O(n)$ or $O(n \log n)$ (and not $O(n^2)$ or more complex),
2. (Storage) the auxiliary data structure should be represented in much smaller space (storage or memory) compared to storing shortest path lengths between all node pairs,
3. (Fidelity) path length (estimation) queries using the auxiliary data-structure should return distances which are as close as possible (if not equal) to the actual distances,
4. (Query Speed) the time required to query the distance between any two nodes should be very small (e.g., small fraction of a second).

In this paper, we focus on a distance oracle for the large- scale graphs that models interactions arising naturally in the real-world, as in online social networks, call

*Bell Labs, Alcatel-Lucent, Dublin, Ireland
†Bell Labs, Alcatel-Lucent, Murray Hill, NJ, USA

graph, co-authorship, citations, hyperlinks in world wide web and similar graphs. We refer to these graphs as real-world graphs. In recent years, considerable effort has been expended in developing approximate distance oracles on these graphs, e.g., see [40, 37, 32, 39, 10, 5, 31], however these heuristics lack a theoretical foundation that would explain their observed accuracy in practice. On the other hand, there are theoretical results [19, 12, 14, 11, 1, 20, 13] which provide guaranteed approximation bounds for specific graph classes. However, the accuracy of many of these methods has not been empirically evaluated on real-world graphs.

Our Contribution. We present a novel distance approximation oracle that leverages the notion of graph hyperbolicity [20] observed in real-world networks [27, 26, 22, 15] and specifically uses the 'hyperbolic core' of the graph [27, 21] for a spanning tree approximation. Hyperbolicity captures the geometric notion of negative curvature in smooth geometry, which we formally define in Section 3 in the context of a graph; intuitively, and crucially, as observed recently [27, 21], it expresses the case that a fixed fraction $\Theta(n^2)$ of all shortest paths traverse a small set of nodes in the graph, thus relative to this small set of nodes, the graph is tree-like in some fundamental ways, a property which we exploit. Our approach also bridges the gap between a) the theoretical understanding of tree approximations for hyperbolic graphs [20, 19, 12, 14, 11, 1, 13] and b) the recent practical distance approximation solutions [40, 37, 32, 39, 10, 5, 31] which exhibit high accuracy.

Our approach is distinct from both sets of prior work since we use neither 1) any kind of hyperbolic embedding [40, 39], nor 2) any form of Gromov-type tree contraction and labeling schemes [12, 19, 14]. We construct a breadth-first search spanning tree prioritized on the nodal betweenness centrality and its approximation by nodal degree, as detailed in Section 4.1. The height of these trees is almost always $O(\log n)$[1] and we use this fact to encode distances in trees with $O(n \log n)$ bits (or $O(n)$ words) to support $O(\log n)$ distance queries. Thus, our proposed oracle 1) creates tree approximations rapidly, 2) uses $O(n)$ words of space for storage for an n-node graph, 3) has high fidelity and low distortion (due to the tree root being in the hyperbolic core) and 4) returns queries very rapidly, meeting the key criteria we set out for an effective oracle. We also demonstrate empirically that the accuracy of the proposed oracle on large real-world graphs is significantly better than the theoretical bounds and competitive with the best practical solutions.

Empirical and Synthetic Datasets. Our empirical comparison is carried out on a wide range of real-world and synthetic graphs, representing a variety of different topological structures from the irregular connectivity of *small world* graphs to the geometric symmetry and regularity of grid and *hypergrid* graphs. Specifically, we experiment on two online social graphs from Facebook, i.e., Santa Barbara and New York (from [38]), two call graphs from anonymous telecom operators, two collaboration networks [24], a Google news graph, a Peer-to-Peer network, a web graph, and two synthetic networks, that we call FlatGrid and HyperGrid. FlatGrid is a square lattice. HyperGrid is a bona fide hyperbolic locally planar graph with degree 7 for interior nodes, lower degrees for boundary nodes, and triangular faces. In Table 1 we summarize the topological and geometric characteristics of these datasets.

Table 1: 9 Real and 2 Synthetic Benchmark Graphs

Dataset	#Nodes	#Edges	Avg. Deg.	Topological Structure
GoogleNews	15.8K	164.1K	21	News
AstroPhysics	18.7K	216.8K	23	Collab.
Facebook SB	26.6K	968K	73	Social
Gnutella	62.6K	210.5K	7	Peer-to-Peer
Call Graph I	631.6K	822.9K	3	Call Graph
BerkStanford	685.2K	7.3M	21	Web
Facebook NY	905.7K	10.6M	23	Social
DBLP	1.1M	4.7M	9	Collab.
Call Graph II	47.2M	329.3M	14	Call Graph
FlatGrid	10.0K	19.8K	4	Grid
HyperGrid	29.3K	51.6K	4	Hyperbolic

Outline of Paper Section 2 summarizes the relevant prior work on distance oracles. In Section 3, we describe the notion of graph hyperbolicity that is at the core of our proposed geometric oracle. Next, in Section 4, we describe the proposed distance approximation oracle and also present various alternative oracles for benchmarking. Subsequently in Section 5, we describe the experimental methodology and summarize our results.

2 Related Work

Theoretical Bounds on General Graphs. There is a rich body of literature on distance oracles for general graphs, including the special case of distance labeling scheme where the distance for query node pairs is estimated by merely using labels associated with the query nodes, and the related problems of graph spanners. The seminal work of Thorup and Zwick [36] described a distance oracle that gives $2k-1$ approximation with $O(k)$ query time, $O(kn^{1+1/k})$ space and $O(kmn^{1/k})$ preprocessing time on an arbitrary weighted undirected graph with n nodes and m edges, for any integer $k \geq 2$. The preprocessing time and the query time of this distance oracle were subsequently improved (c.f. [9, 8, 7, 25]), but the space versus approximation factor trade-off has remained almost the same. In fact, various lower bounds have been proved under plausible conjectures (c.f. [28]

[1] This is due to the well-cited small world property of real-world graphs where the diameter is observed to scale like $O(\log n)$ or even faster (unless the graph has a dominant line in it).

and the references therein) for the space versus worst-case approximation factor trade-off. These lower bound results suggest that it is unlikely that a distance oracle can result in a significantly better trade-off for general weighted graphs.

Empirical Work on Road Networks. In contrast to the theoretical work on general graphs, there has been considerable algorithm engineering and experimentation work on road networks for navigation applications using global positioning systems (GPS). These solutions (e.g., [6, 17, 18, 33]) crucially rely on many specific characteristic properties of road networks such as the existence of small natural cuts, a grid-like structure, highway hierarchies, guiding the search towards the target using the latitude and longitude of the target location, etc. On other graph classes such as those from online social networks, equivalent solutions are not generally known to produce equally good approximations.

Approximate Distances on Real-World Networks.

Distance oracles have been investigated both from theoretical and practical perspectives. As described earlier, our focus is on distance oracles that provide accurate distance estimates on large real-world graphs rapidly (a few microseconds) while having scalable preprocessing and near-linear storage requirement. This category includes a number of recent practical heuristics [39, 40, 10, 37, 32, 31]) that aim to estimate distance by embedding the graph in geometric spaces, like Euclidean or Hyperbolic, or by extracting different kinds of approximating trees from the real graph. Other heuristics, such as, [37, 10] or the landmark based approaches with diverse seeding strategies [30] use variants of breadth first search (BFS) trees. We also remark that the sketch-based distance oracle by Sarna et al. [34], that engineers a distance oracle with provable multiplicative guarantees, is similar in nature to other approaches cited above. On the other side, there are some theoretical approaches [11, 1, 20, 13] that prove worst-case bounds on accuracy for specific graph classes such as those with power-law degree distribution or those with small graph hyperbolicity [12, 19, 14]. These techniques have not been evaluated on large real-world graphs. We evaluate some of these techniques for benchmarking of our oracle's performance.

Exact Distance Oracles. Since exact distance oracles require the computation or storage of all pair shortest path, we do not consider these (e.g., [5, 2] and the references therein) in our comparison. We observe that even with the best combination of engineering insights, all-pairs shortest path computation remains far too slow for large graphs with $\sim 10^7$ nodes or edges and beyond, and this is especially unmanageable when these graphs do not fit in the main memory of the computing device. Also, we do not consider the oracles that have $\Omega(n^\epsilon), \epsilon > 1$ query complexity (e.g., [4, 3, 29]), as these are unlikely to yield efficient solutions.

3 A Geometric Distance Oracle

In this section, we present an overview of graph hyperbolicity and discuss how graph hyperbolicity may be exploited to design an effective distance oracle. Recent studies [27, 15, 22] show that large-scale networks, from IP-layer connectivity, citation, collaboration, co-authorship and friendship graphs, exhibit strong intrinsic hyperbolicity.

δ-Hyperbolicity Intrinsic hyperbolicity, by which is meant hyperbolicity without any embedding of the graph into some Euclidean and other space, captures the geometric notion of negative curvature in smooth geometry. In the past two decades, the notion of negative curvature has been successfully exported to metric spaces, which are more general and less restrictive than Euclidean spaces. This generalized notion of curvature lends itself naturally to the investigation of (large-scale) curvature in graphs [27, 22]. A simple description of δ-hyperbolicity in a (path) metric space is that the three sides of *any* shortest-path triangle X, Y, Z always come within a certain fixed distance δ of each other, where δ is a fixed minimal constant associated with the graph. In other words, the union of two δ-neighborhoods of any two sides of a shortest-path triangle includes the third, as depicted in Figure 1(a).

An easier to compute equivalent condition is as follows. Given a graph $G(V, E)$ and any four nodes $w, x, y, z \in V$, consider the three sums of distances between opposite pairs of nodes. Specifically, let $S := S(w, x, y, z) = d(w, x) + d(y, z) \leq M := M(w, x, y, z) = d(x, y) + d(w, z) \leq L := L(w, x, y, z) = d(x, z) + d(w, y)$ as shown in Figure 1(b). Here, $d(x, y)$ is the shortest path distance between nodes x and y in $G(V, E)$; often we will assume the standard 'hop' metric on G. The hyperbolicity δ of a graph may be defined [20] as follows:

Definition 1. *A graph $G(V, E)$ is said to be δ-hyperbolic for some $\delta \geq 0$ if for every four nodes $w, x, y, z \in V$, $(L - M)/2 \leq \delta$ is always satisfied (i.e., the largest two of the three sums of opposite side distances differ by no more than 2δ).*

We refer to δ as the (4-point) hyperbolic constant of the graph[2]. Notice that any finite graph with diameter Δ is trivially Δ-hyperbolic. Thus this definition is insightful only when δ is considerably smaller than Δ, as argued, for example, in [27, 22]. In Table 2 we list the estimated δ values on our benchmark graphs to show that

[2]Even though δ is used to refer to both the 3-point and 4-point constants, for the same graph these two need not have the same value. From here on we use δ to denote the 4-point constant of the graph.

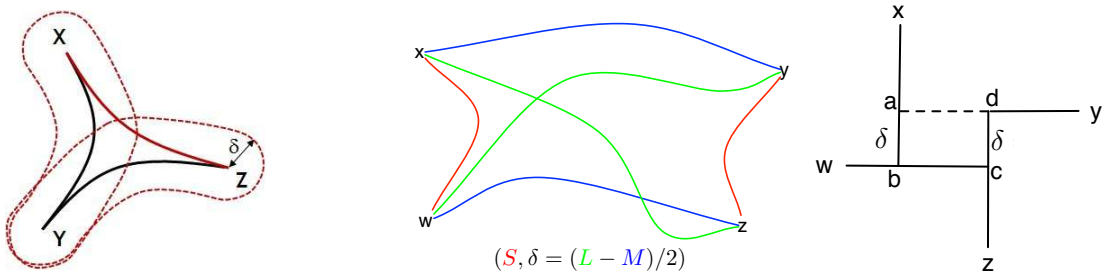

(a) 3-point condition (b) 4-point condition (c) 2δ distortion relative to trees

Figure 1: (a) A schematic representation of the simple-to-visualize 3-point condition. (b) Simpler to compute is the the equivalent 4-point condition, where for any four points w, x, y, z, the largest sum L of opposite pairs ($xz+yw$ in this picture) minus the middle sum M ($=xy+wz$ here) is no more than twice δ. (c) The 4-point condition implies a 2δ deviation from a tree graph, in which the largest among the three sums of opposite pairs of distances of four points is always equal to the medium sum.

Table 2: Key Geometric Characteristics of the 9 Real and 2 Synthetic Benchmark Graphs (based on large samples).

Dataset	Clust Coeff.	max δ	avg δ	Diam
GoogleNews	0.45	1.0	0.03	4
AstroPhysics	0.50	2.0	0.24	~5
Facebook SB	0.22	2.0	0.23	14
Gnutella	0.01	1.5	0.08	~5
Call Graph I	0.14	4.5	0.39	~47
BerkStanford	0.50	2.0	0.16	~5
Facebook NY	0.16	2.0	0.22	~19
DBLP	0.65	2.5	0.24	~23
Call Graph II	0.09	< 4	0.27	~27
FlatGrid	0.00	87.0	7.40	198
HyperGrid	0.33	1.0	0.27	18

indeed on real world networks the δ values are actually very small. For a detailed description of how δ may be estimated for large graphs, see [27, 22, 15].

δ-Approximation of Hyperbolic Graphs with Trees. It has been observed that the distance metric of a δ-hyperbolic graph can be viewed as a δ-approximation of a tree metric, since for any four points $w, x, y, z \in V$ in a tree $T(V, E')$, (for an appropriate set E' of edges) one necessarily has that $S \leq M = L$, thus resulting in $\delta = 0$.[3]

While trees and block graphs are 0-hyperbolic and chordal graphs have $\delta \leq 1$, cycles C_n are $O(n)$-hyperbolic and square grids with n nodes are $O(\sqrt{n})$-hyperbolic.

It turns out that similar to trees, δ-hyperbolic graphs have a non-empty core of nodes whose betweenness centrality is maximal possible, that is, they have nodes whose betweenness centrality scales as $\Theta(n^2)$ (where $n := |V|$ is the size of the graph), as argued in [27, 21]. In other words, δ-hyperbolicity also captures the notion that there is always a non-empty set of vertices, its 'core', where a fraction of *all* shortest paths pass through.[4] This is a crucial property that we shall leverage in our distance approximation oracle.

Theoretical Bounds via Tree Approximations. We focus on three classes of distance oracles that approximate distances on graphs via trees. First, Gromov-like techniques, principally [20, 19, 14], yield an additive guarantee on the distortion on distances via tree approximations of the graph, where the resulting trees are not sub-graphs of the original graph but are typically (based on) contractions. Specifically, Gromov's notion of hyperbolicity leads to approximations of a graph G with trees T that satisfy the following bound: For any two points $x, y \in V(G)$, the shortest path distance $d_{GT}(x, y)$ between them in T has distortion $|d_G(x, y) - d_{GT}(x, y)| \leq 2\delta \log n$. Second, the tree embedding proposed by Abraham *et al* [2]. yields a multiplicative guarantee on the distortion for distances in the tree. A metric space X is defined to be ε-hyperbolic, $\varepsilon \in [0, 1]$, if every set of four points w, x, y, z in X ordered so that $d(w, x) + d(y, z) \leq d(w, y) + d(x, z) \leq d(w, z) + d(x, y)$ satisfies

$$d(w, z) + d(x, y) - d(w, y) - d(x, z) \leq 2\varepsilon \min\{d(w, x), d(y, z)\}.$$

For such a metric space the authors give an algorithm to construct an embedding X into a tree T where for any two points $x, y \in X$ the shortest path distance $d_T(x, y)$ between them in T satisfies $\frac{\max_{x,y}(d_T(x,y)/d_G(x,y))}{\min_{x,y}(d_T(x,y)/d_G(x,y))} \leq (1 + \varepsilon)^{c_1 \log |X|}$. Further, up to the constant c_1, they show that there exists an example where this inequality is tight. Finally, Chepoi and Dragan [12] construct contraction trees (i.e., many nodes of G may map to the same node in T) that approximate the graph distances with an additive distortion not exceeding $\lfloor c(G)/2 \rfloor + 2$

[3] Note that another notion of tree-likeness, tree-width, is an independent property and unlike δ-hyperbolicity its existence in real-world graphs is not common. For instance, snapshots of the internet at autonomous system level have large tree-width, but low hyperbolicity [26].

[4] Note that our usage of the term 'core' here is different from the core of online social networks (that contains most of the graph nodes) and the notion of coresets used in streaming algorithms literature.

where $c(G)$ is the length of the longest chordless cycle in G. This kind of approximation is also appropriate to consider when real-world graphs only have relatively short chordless cycles.

4 Algorithmic Design

In this section, we present an algorithmic framework for designing distance oracles that leverage the small hyperbolicity of real-world graphs. We exploit the hyperbolic core [27, 21] to construct a spanning tree BFS subgraph approximation for the graph, and not tree contractions with non-graph edges as done in the above-referenced three approaches. Since, as cited, for δ-hyperbolic graphs, asymptotically a fixed fraction of all shortest paths in the graph pass through its non-empty hyperbolic core, picking a node from the core as the tree root will give us zero distortion for all node-pair distances whose shortest paths traverse that node. The node with the highest betweenness centrality is thus ideal as the root of the spanning tree: The smaller the hyperbolic constant δ of the graph, the tighter the core, and the larger the fraction of all node-pair paths that traverse a typical node in the core, until for $\delta = 0$ we have a single node to pick as the root to ensure a fixed fraction of all node-pair paths traverse it. The proposed tree oracle therefore starts from nodes with the highest (betweenness) centrality. The tree then expands as a BFS tree by adding unvisited nodes prioritized by their (betweenness) centrality values until all nodes are included in the tree.

Now, since computation of centrality is too expensive on large-scale graphs, we need a good proxy for centrality that is easy to compute. Empirically, it has been observed on real-world graphs that nodal degrees correlate (significantly) well with (betweenness) centrality [27], and Figure 2 shows a typical correlation chart for the FacebookSB network. A degree-based prioritization is then employed both during the selection of initial root of the tree as well as during its expansion.

Four Benchmarked Oracles. In the remainder of this section, we describe three distance oracles that either have proven bounds on distance distortion or are best-of-class empirically. We compare the results with our proposed geometric oracle on the eleven samples of real networks and the two synethetic graphs, as shown in Table 2. More specifically, we examine 1) a representative approach from Gromov-type contraction-based tree approximation with proven bounds: **Gromov Tree**, see Figure 3(c), 2) an oracle based on Steiner trees with proven multiplicative bound: **ε-approximate Steiner Tree**, 3) an empirical best-of-class landmark-based approach that exploits hyperbolic embedding: **Rigel**, and 4) our proposed centrality-based spanning tree oracle: **HyperBFS**, see Figure 3(e), as outlined above.

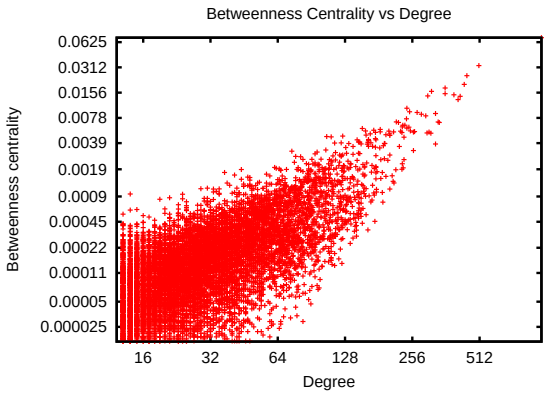

Figure 2: Degree and betweenness centrality of nodes in the Facebook SantaBarbara graph.

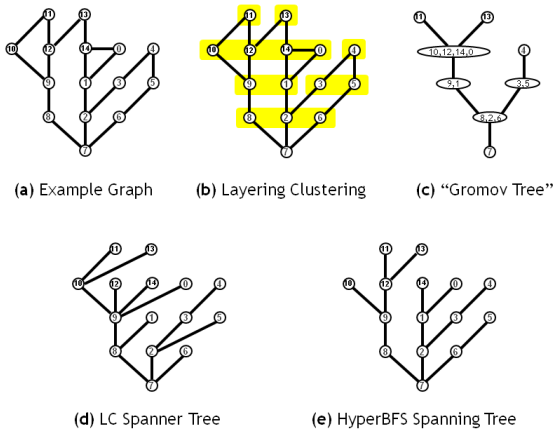

Figure 3: (a) An example graph, (b) its layering partitioning [12], (c) the resulting "Gromov Tree" approximation, (d) the associated tree spanner [14], and (e) our Hyper BFS spanning tree. As it may be seen from this simple example, Hyper BFS is distinct from (c) and (d), while (c) and (d) are very closely associated.

4.1 Hyperbolicity-based Tree Oracle

A δ-hyperbolic graph may be viewed as an approximation to a tree. We thus aim to extract a 'backbone spanning tree' (or a small collection of such trees for better distance approximation) from the input graph to construct our geometric oracle. As stated earlier, we expect a BFS spanning tree with a highly central vertex as root would closely approximate distances in a δ-hyperbolic real- world graph (cf. Section 5 for more details). Also as argued earlier, we may use degree as a proxy for centrality for computational efficiency and select the highest degree node as the root.

Therefore, in constructing the Hyper BFS spanning tree T, we choose the order in which new nodes are added to the BFS tree strictly based on their degree. Algorithm 1 gives pseudocode for constructing the Hyper

BFS tree based on a general vertex ordering π, which in our implementation, and unless stated otherwise, is based on degree.

Algorithm 1 – Hyper BFS with Vertex Ordering

Input: Graph G, vertex ordering π
Output: Tree T

let r be the first vertex in π
set $Q = \{r\}$ and $T = r$.
while Q is not empty:
... **set** v top of Q (remove v from Q)
... **let** $N = N(v)$ ordered by π
... **for each** $n \in N$ not already in T:
..... **push** n on Q
..... add vertex n and edge nv to T
return T

To improve the accuracy of distance estimates, and as practiced previously, we use a small number of different trees, each with a distinct node as root. We construct a collection of such trees rooted at distinct nodes in the hyperbolic core of the graph, which as stated, we approximate by nodal degree. Finally, the distance between two nodes x and y is answered by returning the minimum of the distances between x and y in the (small set of) different trees we construct.

4.2 Contraction Trees with Bounded Approximation

We start with the 'Gromov tree' technique which yields an approximation with guaranteed additive distance distortion based on a contraction tree. Next we outline the approach due to Abraham *et al.* which yields a multiplicative guarantee on the distance distortion in an approximating tree.

Gromov tree approximation. Gromov's notion of hyperbolicity [20] naturally leads to a design for a linear time algorithm for approximating a graph G with a contraction tree T. For any two points $x, y \in V(G)$, the unique shortest path distance $d_{GT}(x, y)$ between them in the Gromov tree T satisfies $|d_G(x, y) - d_{GT}(x, y)| \leq 2\delta \log n$. Following [12], for a connected graph G, we fix a root $r \in V(G)$ and for $i \geq 0$, let $N_i(r)$ be the vertices at distance **exactly** i from r in G, and $B_i(r)$ be the vertices at distance **at most** i from r in G. A Gromov tree T is then obtained by contracting all node pairs $u, v \in N_k(r)$ (for all $k \geq 1$), such that there exists a path between u and v entirely contained outside $B_{k-1}(r)$. For convenience, we define a *i-level connected component* as a maximal subset C of vertices in $N_i(r)$ such that each pair of vertices in C is connected in $G \backslash B_{(i-1)}(r)$. This construction ('layer partition' of a graph), proposed in [12, 14] and also used by [19], determines a tree T satisfying the above properties in linear time by dealing first with those nodes furthest from r; if we think of trees with the root at the top, we can say 'in a bottom-up manner'. In each level, $N_i(r)$, we contract the i-level connected components into a single node. The bottom-up approach yields that in order to find these components we need only consider the edges completely contained in $N_i(r)$ and those between $N_i(r)$ and $N_{i+1}(r)$. It follows that the total time checking i-level connectivity throughout the execution is linear, and therefore, the algorithm also runs in linear time.

Although this Gromov-type tree construction (per [12, 14, 19]) allows arbitrary nodes to be picked as root, in our implementation we pick a high centrality (with degree as proxy) node. We found that this modification considerably improves the accuracy of distance estimates via 'Gromov type' contraction trees. Hereafter, we refer to our modified and improved implementation of prior work (such as [12, 14, 19]) as the *Gromov Tree*. Pseudocode of our implementation is given in Algorithm 2.

Algorithm 2 – Gromov Tree

Input: Graph G
Output: Tree T

let r be high degree node
determine sets $N_0(r), \ldots, N_\ell(r)$
set $T = G$
for $i = \ell$ downto 0:
... **for each** i-level connected component C of $N_i(r)$:
..... contract C in T.
return T

In the construction of 'Gromov tree', the distance $d_{GT}(x, y) \leq d_G(x, y)$. This allows us to use the above technique for a lower bound on the distance estimate and, thereby, compute an approximation range around the real distance (more on this in Section 5.5).

Lemma 1. *For any two nodes x, y in G, $d_{GT}(x, y) \leq d_G(x, y)$*

Proof. A key observation here is that in the construction, we only contract nodes and do not delete any edge. This contraction of nodes may coalesce multiple edges into a single edge. Thus, for each edge $\{u, v\} \in E(G)$, we either have that u and v are contracted to the same node in T or $\{C(u), C(v)\} \in E(T)$, where $C(u)$ is the contracted nodes of u in T. In particular, this means that for $\{u, v\} \in E(G)$, we have $d_{GT}(u, v) := d_{GT}(C(u), C(v)) \leq 1$.

Consider the shortest path $P(x, y)$ in G (with number of edges $|P(x, y)| = d_G(x, y)$). For each edge $\{u, v\}$ in $P(x, y)$, it holds that $d_{GT}(u, v) \leq 1$. Thus, $d_{GT}(x, y) \leq \sum_{\{u,v\} \in P(x,y)} d_{GT}(u, v) \leq \sum_{\{u,v\} \in P(x,y)} 1 = |P(x, y)| = d_G(x, y)$. □

We note in passing that [14] further refines this contraction tree by expanding each partition back to the original number of nodes by connecting these nodes back

to a single node in the lower layer partition, as shown in Figure 3. We shall not consider this "layer partition spanner trees" further as these only add/subtract a fixed amount to the "Gromov Tree" approximation (zero distance for nodes in the same partition change to 2).

ε-approximate Steiner tree. A slightly different version of hyperbolicity was studied by Abraham *et al.* in [1] and it too leads to a tree approximation but with a guaranteed multiplicative distortion. As in the Gromov construction, their algorithm fixes a root $r \in V(G)$ and consider $N_i(r)$. The tree T is then constructed on vertex set $V(G) \cup S$, where S is a set of Steiner points, one for each connected component in the graph induced on each $N_i(r)$, $i > 0$. The edge set is defined as follows. For a connected component C in the graph induced on $N_i(r)$, where s_C is its corresponding Steiner point, we add edges cs_C of weight $\frac{1}{2}$ for each c in C and edge $s_C x$ of weight $\frac{1}{2}$, where x is some vertex of $N_{i-1}(r)$ connected to some vertex of C. Pseudocode is found in Algorithm 3.

Algorithm 3 – ε-approximate Steiner Tree

Input: Graph G
Output: Tree T

let r be high degree node
determine sets $N_0(r), \ldots, N_\ell(r)$
set $T = \emptyset$
for $i = \ell$ downto 0:
... **for each** connected component C of $N_i(r)$:
..... add vertex s_C to T
..... add edge $s_C x$ of weight 1/2 to T (for one $x \in N_{i-1}(r)$ which is adjacent to $c \in C$)
..... **for each** $c \in C$:
..... add vertex c and edge cs_C of weight 1/2 to T
return T

4.3 Embedding in Geometric Space

In this section, we present oracles that embed a graph into some multi-dimensional geometric space. We pick these oracles for special attention since they typically outperform other distance approximation oracles and thus help benchmark the fidelity of our oracle. Hyperbolic embedding involves explicitly mapping the nodes of the graph into points in the hyperbolic space. (We observe here once more that intrinsic hyperbolicity is not the same as hyperbolic embedding used in these distance oracles based on [23, 40].)

Some oracles in this category, such as Orion [39] and that of Qi et al. [31] use the L_2 norm distance function in $\mathbb{R}^{10}$. Others such as Rigel [40] use a *Hyperboloid* model [35] with curvature c, where the distance between two d-dimension points $x = (x_1, x_2, \ldots, x_d)$ and $y = (y_1, y_2, \ldots, y_d)$ is defined as follows:

$$arccosh\left(\sqrt{(1+\sum_{i=1}^{d} x_i^2)(1+\sum_{i=1}^{d} y_i^2)} - \sum_{i=1}^{d} x_i y_i\right) \cdot |c|$$

Algorithm 4 – Rigel & Hyperbolic Embedding of Graph

Input: Graph G
Output: Coordinates for all nodes in $\mathbb{R}^{10}$

Select a set of high degree nodes L
for each $v \in L$:
... BFS(v)
Compute coordinates for $v \in L$ minimizing distortion from each other
for each $v \in V \setminus L$:
... **Select** a subset L_v of L
... $Ord(v)$ = Coordinates that minimize distortion from L_v
return Ord

These oracles embed a graph G by first identifying a small set of nodes with high degrees that are called landmarks. Using BFS from all landmark nodes, they compute the distance of all nodes in the graph from these landmark nodes. Then, a linear program is defined to compute an embedding of these landmark nodes in $\mathbb{R}^{10}$ such that the difference between the actual pairwise distance of these landmarks and their distance estimated using the defined distance function in the embedded space is minimized. Solving this linear program (e.g., via the simplex method) provides the coordinates of landmark nodes.

The remaining nodes in G are then given a coordinate using another linear program. The objective of this linear program is to minimize the difference between the actual distance of the node to a subset L_v of landmark nodes and the distance estimated using the distance function in the embedded space. Once again, the simplex method is used to solve the linear program. Algorithm 4 provides a pseudocode for the Rigel approach.

The approximate distance among each pair of nodes is then their distance in the embedding. Thus, the query time and space per node only depends on the dimensionality of the embedding and for a fixed dimension, it reduces to $O(1)$ (possibly with a high value for the prefactor).

We remark that a plausible reason for the success of these hyperbolic embedding techniques is that in real-world graphs, the set of 'core nodes' of high centrality (as implied by small graph hyperbolicity) and the set of landmark nodes (computed based on degree) includes many of these same nodes as shown in Figure 2.

Note that these coordinate-based systems can both underestimate and overestimate the real distances. Rigel has been shown to be significantly more accurate than high-dimensional Euclidean embeddings [40] and our preliminary experiments also confirmed the same. Hence, we only consider Rigel from this category in our empirical comparison.

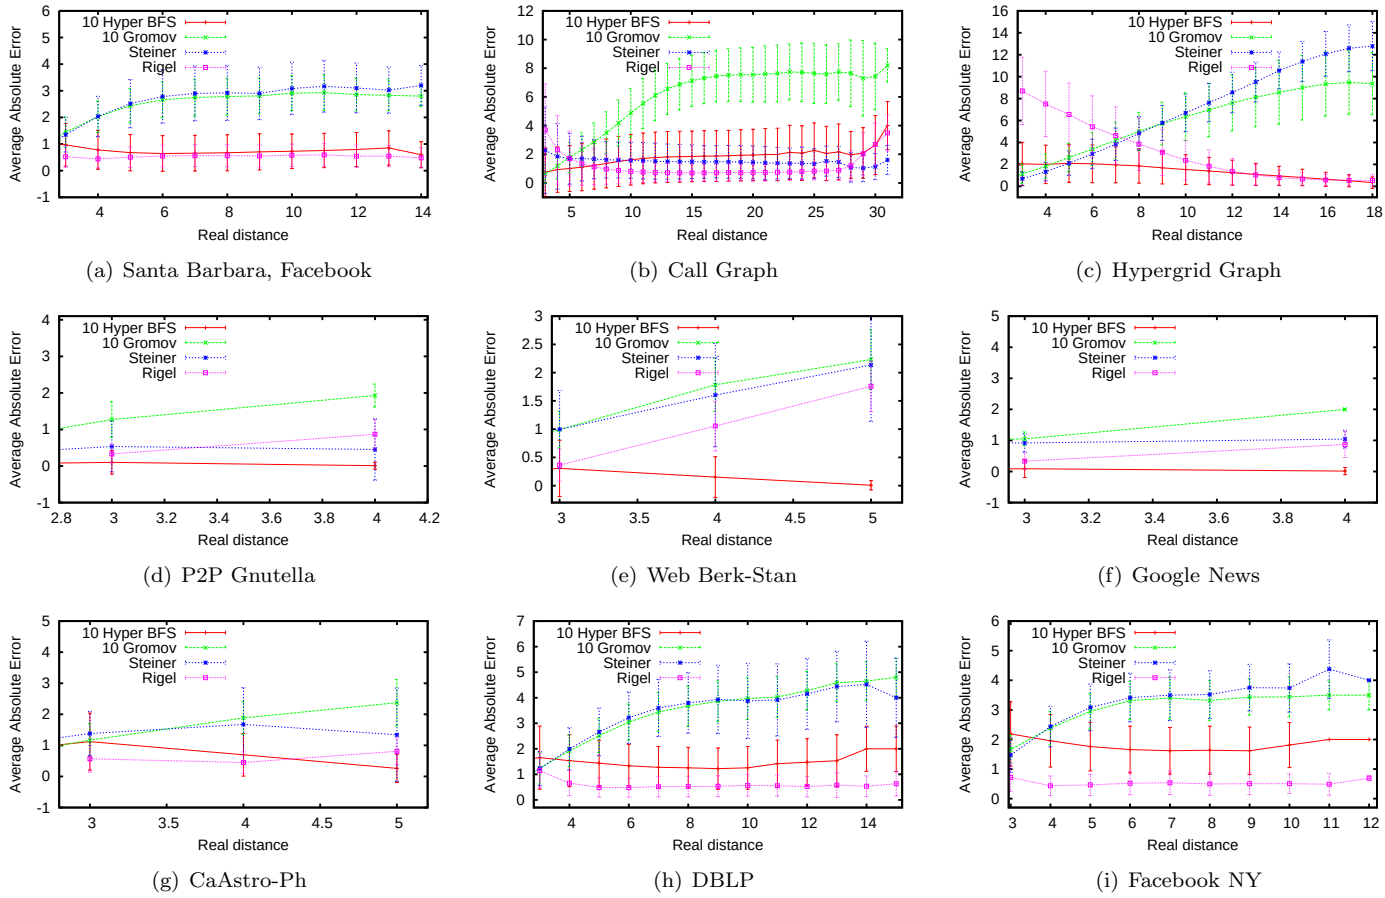

(a) Santa Barbara, Facebook (b) Call Graph (c) Hypergrid Graph

(d) P2P Gnutella (e) Web Berk-Stan (f) Google News

(g) CaAstro-Ph (h) DBLP (i) Facebook NY

Figure 4: Average absolute error of various approximation techniques on various synthetic and real-world graphs.

5 Experimental Evaluation

In this section we describe the methodology used in our experiments, followed by a detailed analysis. The goal of our benchmarking is to determine if the observed δ-hyperbolicity of real-world or synthetic networks yields a low cost and effective distance approximation competitive with or better than the 1) best-of-class or 2) theoretically guaranteed approaches, as discussed.

5.1 Experimental Setup

The networks we study range from 10s of thousands to 100s of millions of nodes and edges. Storing $(|V|(|V|-1))/2$ distances between all node pairs results in Gigabytes of storage for small size graphs and quickly hits hundreds of Terabytes for our larger graphs, which is impractical. To ensure that accuracy is measured over a large enough sample of all node pairs in a graph, we use two different approaches. When the studied graph is small enough, we compute the exact distortion of distance for each pair of nodes. For large graphs we compute the distance distortion by sampling a large enough set of node pairs to ensure that the mean error due to sampling is sufficiently small.

A few observations are in order. First, for each of the Hyper BFS and Gromov Tree techniques, we construct not just one tree but a small collection of trees rooted at nodes in the hyperbolic core of the graph, i.e., those with the highest centrality, approximated with degree centrality for computational efficiency. Once we obtain this collection of trees, the distance between two nodes x and y is answered by returning the minimum of the distances between x and y in the different trees of the Hyper BFS. For the Gromov Tree, the maximum distance gives a better lower bound to the graph distance, since Gromov Trees are contractions. Second, we have verified experimentally that 10 such trees are enough to provide very good distance approximations for the Hyper BFS and our implementation of Gromov Trees. This is an improvement relative to prior work [10], where it was observed experimentally that 20 of their spanning trees were needed for a similar level of accuracy. This improvement is likely due to our selection of roots from the centrality core of the graph. Finally, although we have computed and carefully examined all four oracles on all 11 benchmarked graphs, in what follows we present only the most representative plots.

Measures of Distortion. We use three measures

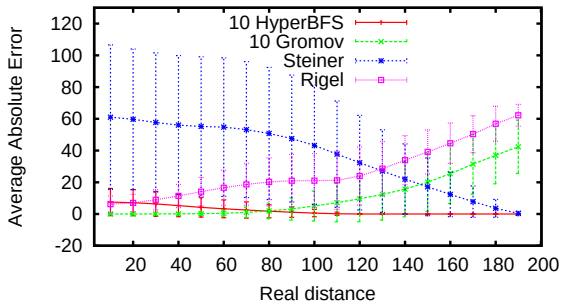

Figure 5: Average absolute error of various approximation techniques on a SquareGrid network

for evaluating the performance of the various distances approximations on each graph. Each of these measures compares the *ground truth* graph distance d_G with the approximate distance given by the proxy d_A and captures a different performance feature.

Definition 2. *Let x, y be vertices of a graph G and let d_A be the distance approximated by a distance oracle.*

- *The* additive distortion *between x and y with respect to d_A is $d_G - d_A$.*
- *The* absolute distortion *between x and y with respect to d_A is $|d_G - d_A|$.*
- *The* multiplicative distortion *between x and y with respect to d_A is $\frac{|d_G - d_A|}{d_G}$.*

The selection of proper distortion metrics is critical in assessing the quality of each distance approximation model since each metric captures distinct aspects of the distortion error and may have different consequences for different applications. The choice of a measure is also important in the context of different kind of bounds known for different approaches. For instance, the Gromov tree approach has additive guarantees while ε-approximate Steiner tree has multiplicative guarantees on distance errors. We visualize each of these distortion metrics by plotting its expected value conditioned on the value of d_G being fixed. This allows us to study the accuracy of various techniques for node pairs at different distances. This is important as applications, like graph segmentation, SVD decomposition, recommendation systems or influential node detection, may require different accuracy guarantees across the range of node distances, i.e. neighbor nodes, diametrically opposite nodes or node pairs with distances in between the two extremes. For instance, recommendation systems may require highly accurate short range distances around a particular user to detect similar users in a particular radius; on the contrary, influential node detection would require accurate long distances to determine the centrality of a particular node.

5.2 Methods' Fidelity

Comparing Oracles Based on Absolute Error. Our first key observation is that the error due to our oracle in estimating shortest path distances is very small. The 10-Hyper BFS approximation results in an absolute error of less than 2 in all instances. Figure 4 shows results on eight real graphs and one synthetic graph, HyperGrid. We next compare the accuracy of the Hyper BFS approach with the best practical oracle system, i.e. Rigel. This system is particularly optimized for networks with Power-Law degree distribution like Santa Barbara(Figure 4(a)) and a Call Graph (Figure 4(b)). However, it significantly loses accuracy for pure hyperbolic graphs like Hypergrid(Figure 4(c)) for short distances. We note that our proposed Hyper BFS technique shows intriguingly low mean error on real-world graphs: on the same range as Rigel on real-world graphs and better for others. Figure 4 also shows that the other techniques, 10 Gromov and Steiner trees, achieve poor results, particularly for node pairs at large distances. Compared to 10-Hyper BFS and Rigel, the Gromov tree contraction technique results in large errors, not only on the call graph and SantaBarbara Facebook graph shown in these figures, but also on the other real- world graphs that we considered. This technique contracts many nodes into the same high degree nodes, thereby underestimating their distances. For instance, on node pairs at a distance of 13 from each other on the call graph, it has an average absolute error of around 6.5, i.e. around 50%.

Short-Long Distance Approximation Accuracy. Here we study how different topological structures impact the fidelity of various techniques on different distance lengths. An important observation that follows from Figure 4, is that Rigel and 10-Hyper BFS result in better accuracy for node pairs with long distances while Gromov performs consistently poorer for such node pairs. We expect that the consistent accuracy of 10-Hyper BFS and Rigel, for farther node pairs, is due to their having nodes with high betweenness centrality (as approximated by degree centrality) at the root or center of their embedding. This ensures that the distances for a large number of node pairs whose shortest path passes through the core nodes are correctly estimated. Interestingly enough, doing the same for 'Gromov Tree' did not seem to improve its accuracy by much, possibly due to its large number of nodal contractions. For short distances instead, Rigel accuracy, in particular on the HyperGrid (Figure 4(c)), leads to a mean absolute error of 8 for the real distance of 3, but gets more accurate for node-pairs at larger distances. This is clearly due to the fact the a planar graph hardly fills a hyperbolic space and hyperbolic embedding is a poor fit. Finally, on the SquareGrid network (Figure 5), which is typical of planar graphs and close to road networks, some techniques like Rigel and

Steiner have very different behaviors compared to the previous networks. Rigel, for instance, estimates long distances with a very large absolute errors. On the contrary, Steiner shows a very high error for distances less than 100 and then for some reason it dramatically improves for very long distances. The fact that these techniques perform well on real-world graphs and synthetic HyperGrid graphs but not on SquareGrid graphs, suggests that the reported success of these techniques most likely relies on the intrinsic hyperbolicity of real world graphs and use of root nodes in the hyperbolic core, two properties not reported heretofore. However, 10-Hyper BFS seems to cope with flatness well, most likely due to its BFS spanning tree structure than its other features such as use of highly central nodes as root, which a square grid does not posses.

Fidelity Versus Theoretical Bounds. Figures 6(a) and 6(b) show the additive and multiplicative errors of the various techniques on SantaBarbara Facebook graph, to provide a quantitative comparison against the respective theoretical bounds provided by the Gromov and the Steiner constructions. For δ-hyperbolic graphs, the theoretical worst-case bound for the Gromov tree based approach is an additive factor of $\delta \log n$. For the social networks we considered, δ was a small constant, but owing to large sizes of the graph, this theoretical bound of $\delta \log n$ is more than the diameter of the graph. However, in Figures 6(a) we observe that the average additive error of Hyper BFS is much smaller than the theoretical bound, and the Gromov tree contraction is evidently the most erroneous compared to other techniques considered. Similarly, the theoretical guarantee for ε-approximate Steiner tree approach is quite large for the studied graphs, in Figure 6(b) we observe a comparatively small multiplicative error of less than 0.6 for most node pairs on this graph. This provides further evidence that the worst-case bounds for the accuracy of these oracles are pessimistic and the average error on real-world graphs is actually much less. Rigel and 10-Hyper BFS approaches are significantly better the theoretical bounds, by exhibiting both a smaller additive and multiplicative errors compared to the Gromov and Steiner techniques. These results confirm that the shortest paths in these graphs are well-captured by 'backbone trees', which is in line with our understanding of the intrinsic hyperbolicity of the real-world graphs.

Exploiting General Hyperbolic Topologies. To reduce the dependence of Hyper BFS on the correlation between degree and betweenness (centrality) of nodes and, thereby, to make this technique more accurate on a wider range of graph classes, we consider alternative strategies for selecting the root nodes for our tree oracle. We identify a node with high closeness[5] centrality by selecting a random node u, finding the node m_u that is at maximum distance from u, selecting a node m_{m_u} at maximum distance from m_u and returning the node in the middle of a path between m_u and m_{m_u} in the graph. We note that similar techniques have been used to approximate diameter of large graphs (see e.g., [16]).

[5]A high closeness central node is a node with low shortest path distances to all other nodes in the graph.

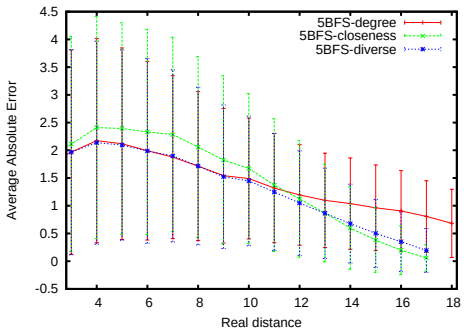

Figure 7: Average absolute error of various Hyper BFS approaches with diverse seeding on the HyperGrid network.

We aim for a technique that would work well on a general graph topologies, particularly real-world social interaction graphs. To this end, we consider a diverse seeding strategy, in which the first seed node is random, the second seed is selected at maximum distance from first seed, the third seed is selected based on closeness centrality (as described before) and the remaining seeds are selected from among high degree nodes. We found that this *diverse* seeding strategy performs well on various graph classes, and in particular on HyperGrid, its accuracy is close to the best results from degree based root selection and closeness based root selection methods (cf. Figure 7).

5.3 Role of Hyperbolic Core

The above benchmarking tests show that Hyper BFS is at least as fast as other tree-based oracles and gives distance estimates comparable to the best. We recall that Hyper BFS constructs a spanning tree of the graph based on the ordering of nodes by their centrality, with a root in the hyperbolic core of the graph. What aspect of Hyper BFS is key to it performance? Here we argue that the selection of the root node of the BFS spanning tree in the hyperbolic core is likely the most critical element in making Hyper BFS a strong distance oracle for real-world networks. We show this by two sets of experiments; The first set keeps the root node in the hyperbolic core but changes the rest of the spanning BFS tree, and the second set changes the root but keeps everything else the same.

Figures 8(a)-8(b) show comparisons on two sets of benchmark real-world networks between Hyper BFS and another implementation of Algorithm 1, which we label Hyper BFS(inc). In this implementation of Hyper BFS,

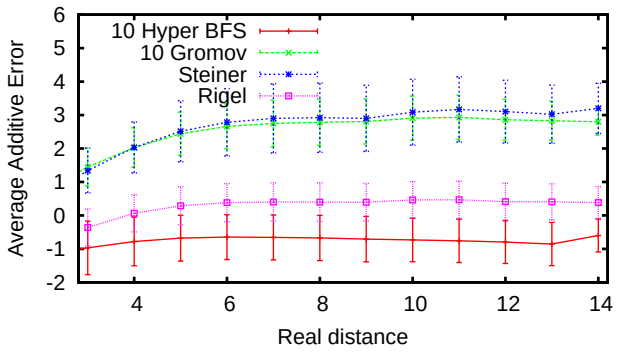

(a) Average Additive Error

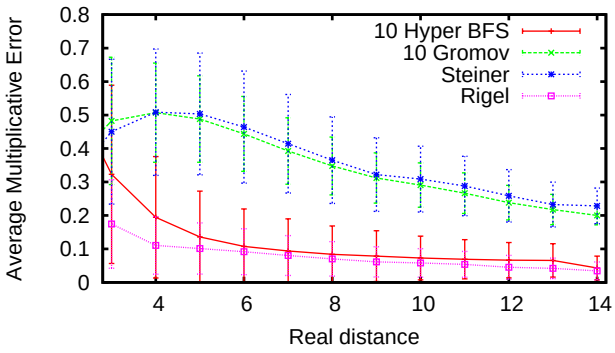

(b) Average Multiplicative Error

Figure 6: Comparing accuracy of the various approximation techniques on the SantaBarbara Facebook graph.

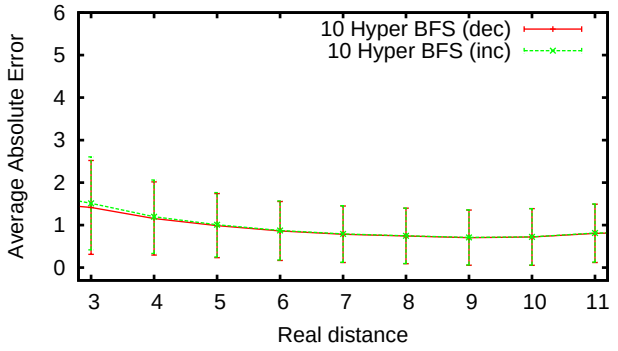

(a) DBLP

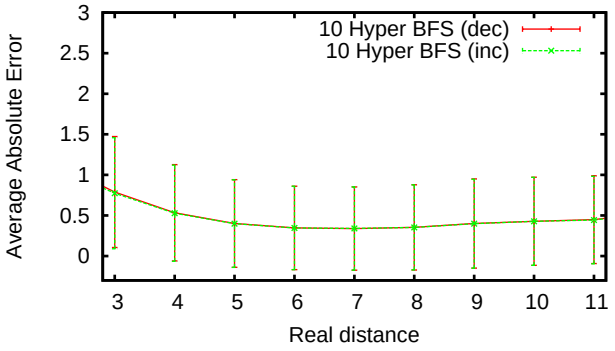

(b) Santa Barbara, Facebook

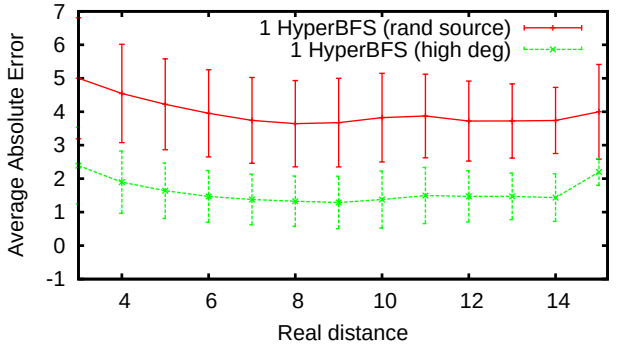

(c) DBLP

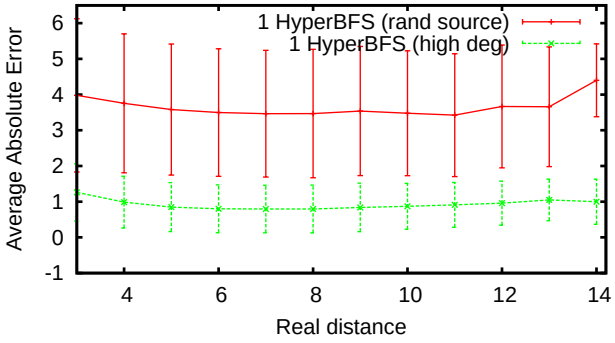

(d) Santa Barbara, Facebook

Figure 8: Top (a)-(b): Comparison of Hyper BFS and Hyper BFS(inc) with increasing nodal degree, on two representative networks both with highest centrality node as BFS root. Bottom (c)-(d): Two instances of Hyper BFS on the same two networks, one with highest centrality node as root (green) and the other with a random node as root (red).

the root node is still selected from the hyperbolic core (i.e., has the highest centrality) but the ordering of the subsequent nodes in the BFS spanning tree is in reverse of Hyper BFS, that is, in order of increasing degrees. As it can be seen, the results are almost identical in terms of average distortion compared to Hyper BFS. However, when we start the Hyper BFS tree with a random root node as shown in Figures 8(c)-8(d), then we see a factor 2-4 increase in the absolute error compared to Hyper BFS where the root has highest centrality. This result, perhaps initially surprising, is consistent with our understanding of the structure of real-world graphs, since the key property of δ-hyperbolic networks is the existence of a core whose centrality is $O(n^2)$, thus ensuring that a large fraction of all shortest paths automatically traverse this small set and for which distance error is zero

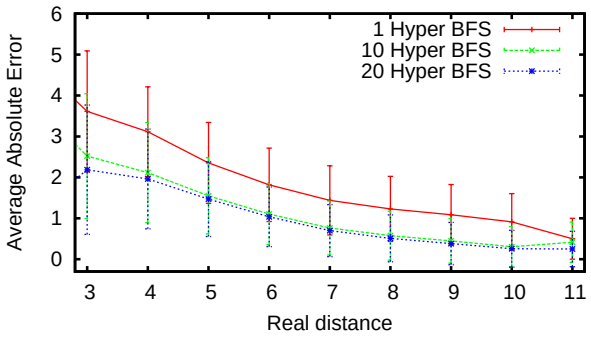

Figure 9: Hyper BFS accuracy on a 50M node call-graph with different number of trees, i.e. from 1 to 20.

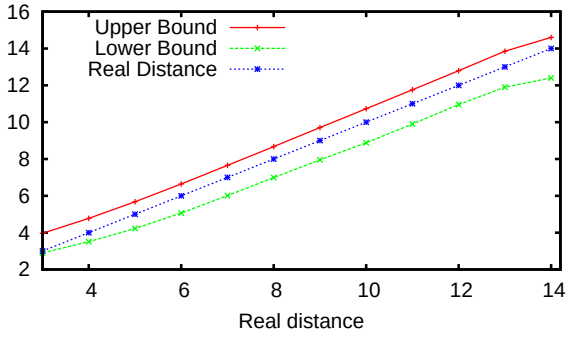

Figure 10: Upper and lower bounds and the width of the approximation range by using 10-Hyper BFS tree and 20 Gromov trees on the SantaBarbara Facebook graph.

Table 3: Computational Time of Hyper BFS on Call Graph II.

Loading Graph	Hyper BFS Tree	1M Queries
250 sec	50 sec	25 sec

for that many node pairs.

5.4 Scalability of Hyper BFS

The Call Graph II has ∼50 million nodes and ∼300 million edges. But as it can be seen from Figure 9 and Table 3, the total run time for the Hyper BFS on this network was under one minutes (under 10 minute for 10 Hyper BFS), including one million node-pair distance queries that were completed in 25 seconds, all on a 2.4 GHz Intel(R) Xeon(R) processor with 190 GB of RAM. We do not know of faster distance approximation techniques with error as small as shown in Figure 9. In our implementation of Hyper BFS, we used a standard tree labeling scheme whereby we store the list of parent nodes of each node to the root, nk indices for k Hyper BFS trees on n nodes. Since the height of the Hyper BFS tree is $O(\log n)$, a query has $O(k \log n)$ complexity. We did not run the other three oracles on this data set as competitive implementations for a graph of this size required substantial optimization of the code and investment of time. We expect Rigel to require several hours for computation of its embedding phase but from there it is likely competitive with Hyper BFS in completing one million queries in 20-30 seconds and based on our observations from other sample real networks, we expect similar or marginally better accuracy.

5.5 Bounding the Hyper BFS Error

Since the Hyper BFS is a spanning tree of the graph, the distance $d_H(x, y)$ in the Hyper BFS tree is an upper bound on the distance $d_G(x, y)$ in the graph for any node pair (x, y). On the other hand, a Gromov tree is a contraction of the input graph and the distance in the Gromov tree $d_{GT}(x, y)$ is a lower bound on the graph distance $d_G(x, y)$ (cf. Lemma 1). Thus, we can use these approaches together to get an approximation range $[l(x, y), u(x, y)]$ with guaranteed upper $(u(x, y))$ and lower $(l(x, y))$ bound for the actual distance $d_G(x, y)$ in the graph, i.e., $l(x, y) \leq d_G(x, y) \leq u(x, y)$. Since $d_G(x, y) \leq d_{GT}(x, y) + 2\delta \log n$, we can use $u(x, y) = \min\{d_H(x, y), d_{GT}(x, y) + 2\delta \log n\}$ and $l(x, y) = d_{GT}(x, y)$. By the above definition of the upper and lower bound, it follows that the width of the approximation range $u(x, y) - l(x, y)$ is less than equal to $2\delta \log n$.

We can reduce the width of this approximation range at the cost of increasing the precomputation time and the storage space by running multiple runs of Hyper BFS tree and the Gromov Tree from different root nodes. Thus, the new lower bound is the maximum over all Gromov Trees. The new upper bound is the minimum over all the hyper BFS trees and the Gromov trees with additive error bound, i.e., $u(x, y) = \min\{\min_H\{d_H(x, y)\}, \min_{GT}\{d_{GT}(x, y) + 2\delta \log n\}\}$. Note that both in the Hyper BFS tree and in the Gromov tree, the distance from the root r to any node $x \in V$ is exact, i.e., $d_H(r, x) = d_{GT}(x, y) = d_G(x, y)$. Thus, if we use n different hyper BFS trees with different root nodes, the minimum over them will yield the exact graph distance. Similarly, if we use n different Gromov trees with different root nodes, the maximum over them will give the exact graph distance, resulting in a approximation range width of zero. However, this extreme point of the solution space requires $O(n^2)$ preprocessing time and $O(n^2)$ storage space.

The accuracy of the Hyper BFS approach is seen to be much smaller than the $2\delta \log n$ theoretical bound for Gromov trees, on the real-world graphs that we tested. This implies that we can obtain a fairly small approximation range by combining the Hyper BFS and the Gromov tree approach. In fact, as Figure 10 shows, the upper and lower bounds on SantaBarbara Facebook graph are quite close and the resultant range with just 10-Hyper BFS trees and 20 Gromov trees is quite small.

6 Conclusion

In this work, we construct a novel distance oracle that is based on the intrinsic hyperbolicity of the graph and specifically leverages its hyperbolic core. Rooting a breadth-first-search spanning tree within the hyperbolic core of the graph ensures that a fixed fraction of all shortest paths have no distance estimation error, and the remaining ones have small error. Furthermore, the smaller the hyperbolic constant δ of the graph, the better such a tree approximation is expected to be.

We tested an implementation of this theory-based framework on a dozen real and synthetic large graphs, from 10s of thousand to 10s of millions of nodes, and found that in practice it leads to surprisingly fast and accurate results, significantly better than anticipated by existing theoretical error bounds that have been proven for alternative tree approximations of the graph. This theoretical framework, especially the existence of the hyperbolic core of the graph, together with evidence from our experiments also suggest that the success of prior heuristic distance oracles may also be due to the same reasons: 1) the underlying hyperbolicity of the real-world graphs and 2) good correlation between nodal degrees and their betweenness centrality in many real-world graphs. In particular, δ-hyperbolicity implies that in these graphs, shortest paths can be well- approximated by appropriately rooted 'backbone spanning trees'.

Another interesting observations is that a simple approach of computing a few BFS trees from high centrality/degree seed nodes provides quite accurate results. To improve this technique further, we considered two strategies: 1) selecting a diverse set of seed nodes for growing BFS trees and 2) to grow the tree by expanding along high degree nodes first. We found that our first strategy helped make the technique more robust – it helped the 20 BFS provide good accuracy even on hypergrid graphs where there was no degree distribution as proxies for betweenness centrality of nodes in the BFS expansion. The second strategy helped on some synthetic graphs, such as hypergrids, but it did not improve distance approximation on real-world networks: it makes a big difference to root the spanning tree approximation at a vertex in the hyperbolic core. We note additionally that compared to the best performing distance approximation oracles, such as Rigel, our approach has a very low computing cost for creating the oracle and thus lends itself well to settings where the graph connectivity changes frequently, such as in large scale dynamic graphs.

We contend that further theoretical understanding of the geometry of δ-hyperbolic graphs can lead to even better distance oracles for real-world graphs. For example, it would be helpful to know how large the hyperbolic core of a δ-hyperbolic graph can be. Another interesting open problem is to find alternative proxies for betweenness centrality in real-world graphs which, as with nodal degrees, is easy to compute.

Acknowledgment. The work of Iraj Saniee and Sean Kennedy was supported by the AFOSR grant no. FA9550-11-1-0278 and the NIST grant no. 60NANB10D128.

References

[1] I. Abraham, M. Balakrishnan, F. Kuhn, D. Malkhi, V. Ramasubramanian, and K. Talwar. Reconstructing approximate tree metrics. In *Proc. of PODC*, pages 43–52, New York, NY, USA, 2007. ACM.

[2] I. Abraham, D. Delling, A. V. Goldberg, and R. F. F. Werneck. Hierarchical hub labelings for shortest paths. In *Proceedings of the 20th Annual European Symposium on Algorithms ESA*, volume 7501 of *Lecture Notes in Computer Science*, pages 24–35. Springer, 2012.

[3] R. Agarwal and P. B. Godfrey. Brief announcement: a simple stretch 2 distance oracle. In *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 110–112. ACM, 2013.

[4] R. Agarwal and P. B. Godfrey. Distance oracles for stretch less than 2. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 526–538. SIAM, 2013.

[5] T. Akiba, Y. Iwata, and Y. Yoshida. Fast exact shortest-path distance queries on large networks by pruned landmark labeling. In *Proceedings of the ACM SIGMOD International Conference on Management of Data SIGMOD*, pages 349–360. ACM, 2013.

[6] H. Bast, S. Funke, and D. Matijevic. Ultrafast shortest-path queries via transit nodes. In C. Demetrescu, A. V. Goldberg, and D. S. Johnson, editors, *The shortest path problem : ninth DIMACS implemenation challenge*, volume 74 of *DIMACS Series on Disrecte Mathematics and Theoretical Computer Science*, pages 175–192. AMS, 2009.

[7] S. Baswana, A. Gaur, S. Sen, and J. Upadhyay. Distance Oracles for Unweighted Graphs: Breaking the Quadratic Barrier with Constant Additive Error. In *35th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 5125 of *Lecture Notes in Computer Science*, pages 609–621. Springer, 2008.

[8] S. Baswana and T. Kavitha. Faster Algorithms for Approximate Distance Oracles and All-Pairs Small Stretch Paths. In *47th Annual IEEE Symposium on Foundations of Computer Science FOCS*, pages 591–602. IEEE Computer Society, 2006.

[9] S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in expected $O(n^2)$ time. *ACM Transactions on Algorithms*, 2(4):557–577, 2006.

[10] L. Cao, X. Zhao, H. Zheng, and B. Y. Zhao. Atlas: Approximating Shortest Paths in Social Graphs. Technical Report 2011-09, University California Santa Barbara, October 2011.

[11] W. Chen, C. Sommer, S.-H. Teng, and Y. Wang. A compact routing scheme and approximate distance oracle for power-law graphs. *ACM Transactions on Algorithms*, 9(1):4, 2012.

[12] V. Chepoi and F. Dragan. A Note on Distance Approximating Trees in Graphs. *European Journal of Combinatorics*, 21(6):761–766, 2000.

[13] V. Chepoi, F. Dragan, B. Estellon, M. Habib, and Y. Vaxès. Diameters, centers, and approximating trees of δ-hyperbolic geodesic spaces and graphs. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, SCG '08, pages 59–68, New York, NY, USA, 2008. ACM.

[14] V. Chepoi, F. Dragan, B. Estellon, M. Habib, Y. Vaxès, and Y. Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmics*, 62(3-4):713–732, 2012.

[15] N. Cohen, D. Coudert, and A. Lancin. Exact and approximate algorithms for computing the hyperbolicity of large-scale graphs. *Publications of INRIA, hal-00735481–version 4, Sophia-Antiplos*, 2013.

[16] P. Crescenzi, R. Grossi, C. Imbrenda, L. Lanzi, and A. Marino. Finding the diameter in real-world graphs - experimentally turning a lower bound into an upper bound. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA*, volume 6346 of *Lecture Notes in Computer Science*, pages 302–313. Springer, 2010.

[17] D. Delling, A. V. Goldberg, I. Razenshteyn, and R. F. F. Werneck. Graph Partitioning with Natural Cuts. In *25th IEEE International Symposium on Parallel and Distributed Processing (IPDPS)*, pages 1135–1146. IEEE, 2011.

[18] D. Delling, A. V. Goldberg, and R. F. F. Werneck. Shortest Paths in Road Networks: From Practice to Theory and Back. *Information Technology*, 53(6):294–301, 2011.

[19] C. Gavoille and O. Ly. Distance labeling in hyperbolic graphs. *Algorithms and Computing*, 3727:1071–1079, 2005.

[20] M. Gromov. Hyperbolic Groups. In S. M. Gersten, editor, *Essays in Group Theory*, number 8 in Math. Sci. Res. Inst. Pub., pages 75–265. Springer-Verlag, 1987.

[21] E. A. Jonckheere, M. Lou, F. Bonahon, and Y. Baryshnikov. Euclidean versus hyperbolic congestion in idealized versus experimental networks. *Internet Mathematics*, 7(1), 2011.

[22] W. S. Kennedy, O. Narayan, and I. Saniee. On the Hyperbolicity of Social and Other Large-Scale Networks. http://arxiv.org/pdf/1307.0031.pdf, 2013.

[23] R. Kleinberg. Geographic routing using hyperbolic space. In *26th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1902–1909. IEEE, 2007.

[24] Lasagne: Laboratory of algorithms, models, and analysis of graphs and networks. `http://amici.dsi.unifi.it/lasagne/?page_id=19`.

[25] M. Mendel and A. Naor. Ramsey partitions and proximity data structures. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 109–118. IEEE Computer Society, 2006.

[26] F. d. Montgolfier, M. Soto, and L. Viennot. Treewidth and hyperbolicity of the internet. In *Proceedings of the 2011 IEEE 10th International Symposium on Network Computing and Applications (NCA)*, pages 25–32. IEEE Computer Society, 2011.

[27] O. Narayan and I. Saniee. Large-scale curvature of networks. *Physical Review E*, 84(066108), 2011. 2009 arXiv.org version.

[28] M. Patrascu and L. Roditty. Distance Oracles beyond the Thorup-Zwick Bound. In *51th Annual IEEE Symposium on Foundations of Computer Science, (FOCS)*, pages 815–823. IEEE Computer Society, 2010.

[29] E. Porat and L. Roditty. Preprocess, set, query! In *Proceedings of the 19th Annual European Symposium on Algorithms (ESA)*, volume 6942 of *Lecture Notes in Computer Science*, pages 603–614. Springer, 2011.

[30] M. Potamias, F. Bonchi, C. Castillo, and A. Gionis. Fast shortest path distance estimation in large networks. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management (CIKM)*, pages 867–876, 2009.

[31] Z. Qi, Y. Xiao, B. Shao, and H. Wang. Toward a distance oracle for billion-node graphs. *Proceedings of the VLDB endowment (PVLDB)*, 7(1):61–72, 2013.

[32] M. Qiao, H. Cheng, and J. X. Yu. Querying Shortest Path Distance with Bounded Errors in Large Graphs. In *23rd International Conference on Scientific and Statistical Database Management (SSDBM)*, volume 6809 of *Lecture Notes in Computer Science*, pages 255–273. Springer, 2011.

[33] P. Sanders and D. Schultes. Engineering highway hierarchies. *ACM Journal of Experimental Algorithmics*, 17(1), 2012.

[34] A. D. Sarna, S. Gollapudi, M. Najork, and R. Panigrahy. A Sketch-based Distance Oracle for Web-Scale Graphs. In *Proceedings of the Third International Conference on Web Search and Web Data Mining (WSDM)*, pages 401–410, 2010.

[35] Y. Shavitt and T. Tankel. Hyperbolic embedding of internet graph for distance estimation and overlay construction. *IEEE/ACM Trans. Netw.*, 16(1):25–36, 2008.

[36] M. Thorup and U. Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.

[37] K. Tretyakov, A. Armas-Cervantes, L. García-Bañuelos, J. Vilo, and M. Dumas. Fast fully dynamic landmark-based estimation of shortest path distances in very large graphs. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management (CIKM)*, pages 1785–1794, 2011.

[38] C. Wilson, A. Sala, K. P. N. Puttaswamy, and B. Y. Zhao. Beyond Social Graphs: User Interactions in Online Social Networks and their Implications. *ACM Trans. Web*, pages 17:1–17:31, 2012.

[39] X. Zhao, A. Sala, C. Wilson, H. Zheng, and B. Y. Zhao. Orion: shortest path estimation for large social graphs. In *Proc. of WOSN*. USENIX Association, 2010.

[40] X. Zhao, A. Sala, H. Zheng, and B. Y. Zhao. Efficient Shortest Paths on Massive Social Graphs, October 2011. (Invited Paper) Proceedings of 7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom).